



# Temporal True and Surrogate Fitness Landscape Analysis for Expensive Bi-Objective Optimisation

Cedric J. Rodriguez  
c.j.rodriguez@lumc.nl

Leiden University Medical Center  
Leiden, The Netherlands

Tanja Alderliesten  
t.alderliesten@lumc.nl

Leiden University Medical Center  
Leiden, The Netherlands

Sarah L. Thomson  
s.thomson4@napier.ac.uk

Edinburgh Napier University  
Edinburgh, Scotland

Peter A.N. Bosman  
peter.bosman@cwil.nl

Centrum Wiskunde & Informatica  
Amsterdam, The Netherlands

## ABSTRACT

Many real-world problems have expensive-to-compute fitness functions and are multi-objective in nature. Surrogate-assisted evolutionary algorithms are often used to tackle such problems. Despite this, literature about analysing the fitness landscapes induced by surrogate models is limited, and even non-existent for multi-objective problems. This study addresses this critical gap by comparing landscapes of the true fitness function with those of surrogate models for multi-objective functions. Moreover, it does so temporally by examining landscape features at different points in time during optimisation, in the vicinity of the population at that point in time. We consider the BBOB bi-objective benchmark functions in our experiments. The results of the fitness landscape analysis reveals significant differences between true and surrogate features at different time points during optimisation. Despite these differences, the true and surrogate landscape features still show high correlations between each other. Furthermore, this study identifies which landscape features are related to search and demonstrates that both surrogate and true landscape features are capable of predicting algorithm performance. These findings indicate that temporal analysis of the landscape features may help to facilitate the design of surrogate switching approaches to improve performance in multi-objective optimisation.

## CCS CONCEPTS

• **Computing methodologies** → **Continuous space search.**

## KEYWORDS

Fitness landscape analysis, surrogate models, expensive optimisation, bi-objective problems

## ACM Reference Format:

Cedric J. Rodriguez, Sarah L. Thomson, Tanja Alderliesten, and Peter A.N. Bosman. 2024. Temporal True and Surrogate Fitness Landscape Analysis for

Expensive Bi-Objective Optimisation. In *Genetic and Evolutionary Computation Conference (GECCO '24)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3638529.3654125>

## 1 INTRODUCTION

In real-world optimisation problems, fitness evaluations can be computationally expensive — ranging from minutes to hours [27] — and multi-objective (MO) in nature. This greatly limits the total amount of possible fitness evaluations to typically a few hundred or thousand. For such expensive-to-evaluate problems, surrogate-assisted evolutionary algorithms (SA-EAs) have been introduced [17]. SA-EAs reduce the required (true) expensive problem evaluations by substituting the true fitness function during parts of the search, with a quick-to-evaluate model (surrogate) of the problem built on a set of true evaluated solutions. In one scenario that is akin to Bayesian optimisation, the SA-EA iteratively optimises the surrogate function, selects one or more surrogate optimised solutions and evaluates using the true function, and updates the surrogate model based on the new set of true evaluated solutions. Because this surrogate is built on a very limited number of solutions, there are likely discrepancies between the true fitness landscape and the surrogate fitness landscape. Therefore, the performance of the SA-EAs is not only dependent on the true fitness function but also on the surrogate fitness function. In this work, we analyse these potential discrepancies between the true and surrogate landscape to better understand the relation between the optimisation performance, the surrogate, and the true fitness landscape during the optimisation.

To properly study potential discrepancies between the true and surrogate fitness landscapes, we need to characterise fitness landscapes during different points of the optimisation because the surrogate landscape may constantly shift due to the surrogate model being updated. Fitness landscape analysis (FLA) can do that, providing insight about the interplay between optimisation algorithms and different fitness landscape. The nature of FLA depends on whether the search space is discrete or continuous, and on how many objectives are formulated. We consider in this work continuous bi-objective problems, and compute features of the true fitness landscape and also surrogate landscapes. In the continuous optimisation domain, many FLA works have focused on algorithm-agnostic (static) approaches, such as Latin hypercube sampling (LHS) [4, 29, 31, 32] and random walks [18, 23, 24]. The nature of the sampling means that the computed landscape metrics may not be closely related to the



This work is licensed under a Creative Commons Attribution International 4.0 License. *GECCO '24, July 14–18, 2024, Melbourne, VIC, Australia*  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0494-9/24/07.  
<https://doi.org/10.1145/3638529.3654125>

landscape that is observed through fitness evaluations performed during an actual optimisation run. In a survey on landscape analysis, Malan and Engelbrecht referred to this as *search independence* [30]. There are also several works which use search-dependent approaches: one of these is continuous-space local optima networks (LONs) [1, 38]. LONs, although valuable, amalgamate information from across the evolutionary timespan into a single mathematical object (the LON). There is also the multi-objective equivalent, Pareto LONs (PLONs) [28]. PLONs have several associated features which have been linked to search performance; however, we do not use them here because they do not fully describe the landscape as observed by an algorithm during optimisation.

Different to the static FLAs, Jankovic and Doerr considered the evolution of exploratory landscape analysis features in a single-objective context [19], finding that static and temporal landscapes do not match. Several works have studied the notion of *trajectory features* (for single-objective problems): these allow an algorithm to begin optimising while logging its trajectory, then computing features from it to understand the problem or make algorithmic decisions [6, 20–22, 40]. These contributions are related to ours because they leverage intermediate landscape information yet focus on single-objective problems. We agree with the view that studying temporal features associated with using an optimisation algorithm, or considering the evolution of them over time may be valuable. Very recently, Alsouly *et al.* studied the evolution of landscape features over time for a set of constrained multi-objective problems [2], finding that the landscape shifts across the course of evolution, and that predicting algorithm performance is possible with temporal features. In that work, several of the features related to the constrained nature of the problems. Two previous studies [33, 34] have computed metrics at different points in evolution for a surrogate-guided evolutionary algorithm, but these were joined together for analysis and the change over time was not separated out and studied. In that case, the sample points used for landscape analysis were based on the true fitness function, and the problems were single-objective; the authors found that several metrics are linked to surrogate-assisted algorithm performance.

Aside from the mentioned works, there are very few articles considering surrogate fitness landscapes. Werth *et al.* explore this direction [41], comparing surrogate landscapes with those of the true fitness function for a small number of single-objective problems. Harrison *et al.* analyse the surrogate landscape for single-objective parameter configuration [16], but do not compare with the true landscape. Generally, there is a lack of literature analysing surrogate landscapes. In the context of multi-objective optimisation, the literature for this — to the best of our knowledge — non-existent.

The present work is the first to conduct analysis of multi-objective surrogate landscapes. Additionally, we compare temporal landscape features of the true fitness landscape and the constantly-shifting surrogate fitness landscape, emphasizing the locally encountered landscape and how it changes over time. Although there has been an initial work studying multi-objective features across evolution [2], this was for constrained problems and did not consider surrogates. Additionally, we consider here a different set of functions and features and argue that this is an under-explored avenue. Another contribution is that we demonstrate the potential associated with building algorithm performance prediction models using features

of *both* the surrogate and true landscape. The rest of this paper is structured as follows: Section 2 will introduce the surrogate models, EAs, and FLA methods, and dimensionality reduction methods that will be considered in this work. Sections 3 and 4 will describe our main methodology and experimental setup, respectively. In Section 5, the experimental results will be presented. Finally, Section 6 consists of the conclusions of the work.

## 2 BACKGROUND

### 2.1 MO Surrogate-assisted EA

Figure 1(a), describes the main process of the type of MO SA-EA that we consider in this work. During the initialization phase, the SA-EA samples a set of solutions using LHS, evaluates them using the true fitness function, and builds the initial surrogate model. After the initialization, the SA-EA is composed of an inner cycle and an outer cycle. The inner cycle consists of typical generational processes of an EA. This cycle initializes with a population containing all previously evaluated solutions using the true fitness function. After the application of the selection and variational operator, offspring solutions are evaluated using the quick-to-evaluate surrogate model instead of the expensive-to-evaluate problem. The cycle terminates as soon as the maximum number of surrogate evaluations have been performed. After the termination of the generational cycle, the surrogate optimisation cycle starts. All surrogate-evaluated solutions are then considered as potential candidate solutions for evaluation with the true fitness function. In the multi-objective context considered for this work, the selection process consists of calculating the domination ranks for all solutions based on their surrogate fitness values. A pre-selection is made by excluding 50% of the solutions with the largest ranks. Thereafter, a set of solutions is randomly selected from the pre-selection for the true evaluation. The true evaluated solutions are then logged and, subsequently, a surrogate model is updated based on this new set of solutions. This SA-EA is described more formally in Algorithm 1.

---

#### Algorithm 1 MO SA-EA

---

**Input:**  $N$  = Number of initial solutions;  $FE_{max}^{true}$  = maximum number of true fitness evaluations;  $FE_{max}^{surrogate}$  = maximum number of surrogate evaluations in each surrogate optimisation cycle;  $\mu$  = number of new selected solutions after each surrogate optimisation cycle;

**Output:**  $A$  = Archive of all non-dominated solutions

- 1: Start first surrogate optimisation cycle and set true function evaluation  $FE^{true}$  to 0
  - 2: Initialize population  $P$  of size  $N$  using random sampling and evaluate solutions using the true fitness function
  - 3: Add initial population  $P$  to archive  $A$
  - 4: **while**  $FE < FE_{max}^{true}$  **do**
  - 5: Build surrogate model  $M$  based on  $A$
  - 6: Optimize  $M$  using EA with  $FE_{max}^{surrogate}$
  - 7: Select  $\mu$  solutions from all surrogate evaluated solutions to evaluate with true fitness function ( $FE^{true} = FE^{true} + \mu$ )
  - 8: Add solutions to the archive
  - 9: **end while**
  - 10: Return  $A$
-

## 2.2 Surrogate models

**2.2.1 Exact interpolators.** In previous expensive multi-objective optimisation algorithms, several point interpolation methods have been considered, such as inverse distance weighting (IDW) [36]. *IDW* is one of the most straightforward and widely adopted [26] point interpolation methods, defined as follows:

$$\hat{y}^{idw}(x) = \sum_{i=1}^n w_i(x) y_i \quad (1)$$

where  $\hat{y}^{idw}$  is the surrogate fitness value of a solution with an unknown true fitness value,  $y_i$  is the true fitness value of the  $i$ th solution in an archive  $A_i$  containing  $n$  solutions, and

$$w_i(x) = \frac{d(x, A_i)^{-1}}{\sum_{j=1}^n d(x, A_j)^{-1}} \quad (2)$$

where  $d$  is the Euclidean distance and  $x$  the unknown solution and the  $i$ th solution. A downside of using *IDW* is that all surrogate values of the solutions outside the parameter bounds of the known solutions approach the average fitness of the known solutions. To overcome this limitation, inverse distance weighting regression (*IDWR*) has recently been developed [12] where the surrogate fitness values of solutions outside the parameter bounds follow the trend line of the fitness value of all known solutions. The surrogate value using the *IDWR* interpolation method is computed as follows

$$\hat{y}^{idwr}(x) = \hat{y}^{idw}(x) + n \frac{\sum_{i=1}^n y_i - n \hat{y}^{idw}(x)}{n^2 - \sum_{i=1}^n d(x, A_i) \sum_{i=1}^n d(x, A_i)^{-1}} \quad (3)$$

**2.2.2 Linear regression.** Another surrogate model is linear regression. To preserve the local structure of the fitness landscape, the surrogate is built on the  $k$ -nearest neighbours (KNN) of the solution with unknown true fitness value. Linear regression is defined as:

$$\hat{y}^{lr-knn}(x) = \beta_0 + \sum_{q=1}^D \beta_q x_q \quad (4)$$

where  $\hat{y}^{lr-knn}(x)$  is the surrogate fitness value,  $\beta_q$  are the coefficient (slope) in each problem dimension,  $D$  is the number of problem dimensions and  $x_q$  are the decision variables.

## 2.3 Reference vector guided EA

The optimizer for the SA-EA in this work is a reference vector guided evolutionary algorithm (RVEA) [7]. The initial population is generated by uniformly sampling solutions in the initialization range of the multi-objective problem. As long as the maximum number of function evaluations is not exceeded, simulated binary crossover [9] followed by polynomial mutation [10] is carried out for generating an offspring population from the parent population. This method for generating offspring solutions is equivalent to other multi-objective EAs such as NSGA-III [11]. After combining the parent and offspring population, the parent population for the next generation is selected using a reference vector guided procedure called the angle penalized distance.

**2.3.1 Selection using angle penalized distance.** Angle penalized distance [7] consists of four mechanisms: generation of reference vectors, assignment of each solution to one of the reference vectors, selection of a solution per reference vector, and the adaption of reference vectors. With RVEA, the objective space is divided into

discrete partitions using reference vectors. The initial generation of these uniform reference vectors is done by firstly generating a set of uniformly distributed reference points. In the next step, each solution in the population is assigned to the reference vector with the smallest angle between the objective vector (the vector between the objective values and the origin) and the reference vector. This partitions the population into several subpopulations, where each subpopulation is associated with a particular reference vector. In the third step, one solution is selected per subpopulation. The selection criterion consists of two subcriteria. The first subcriterion (convergence) is measuring the magnitude of the objective vector, where a smaller objective vector means a better solution. The second subcriterion (divergence) is the angle between the objective vector and the reference vector, where a larger angle means more divergence. These two subcriteria are combined into a single criterion: the angle penalized distance, where the parameter  $\alpha$  specifies the balance between the subcriteria.

## 3 METHODOLOGY

The main processes in the methodology are presented in Figure 1(b). The methodology is composed of a static FLA and a temporal FLA. For the static FLA,  $200D$  solutions are sampled using LHS, where  $D$  is the number of problem dimensions. These solutions are used to extract the true fitness landscape features. In the temporal FLA, for every problem and surrogate model pair, we independently run an optimisation process using an SA-EA. Independent optimisation runs driven by each surrogate model are required because different models might guide the EA to different parts of the search space and encounter different fitness landscapes. During each optimisation, solutions are logged alongside both the surrogate fitness and also the true fitness. From these samples, landscape features and optimisation metrics are extracted for the final analyses.

### 3.1 Sampling during optimisation

During 15 repeats of the optimisation, the locally encountered true and surrogate fitness landscape is sampled at each time point in the following way: from the current population, we repeatedly apply the selection and variational operators until 2000 evaluations have been reached; we thereby sample the vicinity around the population. From this sample, duplicates are removed, meaning that the sample can be slightly less than 2000. Two different sets of solutions are thus logged after every surrogate optimisation cycle:

- (1) Sample of solutions that have been evaluated by the SA-EA within the current surrogate optimisation cycle. The solutions are also evaluated using the true fitness function to acquire the true fitness values.
- (2) The true evaluated solutions for the optimisation.

### 3.2 Feature extraction process

In the feature extraction process, the true and surrogate landscape features are extracted using the FLA method from Liefvooghe et al. [29] designed for continuous multi-objective problems. From the 49 features, we removed any features which displayed missing values in any of the samples. Missing values can occur for some of the features depending on the nature of the sample; for example, if there is only one non-dominated point in a sample, then there




Figure 1: A) Visualisation of the main process of the type of SA-EA that we consider in this work where the (orange) inner cycle represents a generational cycle of an EA with a selection and variational operator. Within this cycle, the solutions are evaluated using a surrogate model. After termination of the generational cycles, the (blue) outer cycle starts, where potentially good surrogate evaluated solutions are selected and evaluated with the true fitness function. B) The general methodology is composed of a static FLA and a temporal FLA where the static FLA utilizes a static sampling strategy while the temporal FLA logs the solutions during the optimisation. A feature extraction process then extracts the true and surrogate landscape features. Furthermore, the performance metric is calculated, which is used for algorithm performance prediction.

will be missing values for the metrics relating to distances between non-dominated points. Additionally, we remove a group of features which are associated with hypervolume; because hypervolume is the measurement we will be predicting, it seemed appropriate to exclude them. The resulting set of 28 features is provided in Table 1 of the supplemental material. The surrogate features are expected to change temporally as the surrogate model is updated after every surrogate optimisation cycle — and indeed, the population is undergoing evolution and the algorithm is moving through different search space regions. Next to the surrogate landscape features, the true landscape features are extracted after every surrogate optimisation cycle. Subsequently, all solutions evaluated with the true fitness function, up to the current surrogate optimisation cycle, are used to quantify the optimisation convergence of the optimisation run. Optimisation convergence is evaluated by first normalising the objectives using the ideal and Nadir points, as recommended in the literature [5]. The Nadir point then corresponds to  $[1, 1]$ . In the context of multi-objective optimisation, we are only interested in regions where there is a trade-off between objectives, which are the solutions between the ideal and Nadir point. The median landscape feature and performance metric is calculated over 15 repeats.

### 3.3 Data representation for analyses

For every problem and surrogate model pair, data is composed of median surrogate landscape features, true landscape features, and a performance metric as a time-series. For computational expense purposes, only the data at the time instances (true function evaluations) 256, 1280, 1792, and 8192 are considered.

## 4 EXPERIMENTAL SETUP

### 4.1 Benchmark problems

The bi-objective BBOB functions (BBOB-BIOBJ) [5] with continuous problem variables from the COCO framework [14] are considered. The 55 problems are a pair-wise combination of 10 well-understood single objective BBOB functions. In this work, we are only considering the first instance of each of the 55 BBOB-BIOBJ problems. We focus on the 20-dimensional instances only. All algorithms have their populations initialized in the default COCO ranges:  $-100$  to

100. The benchmark problems are simulated to be computationally expensive by setting the maximum number of true function evaluations to 8192. The maximum runtime of each optimisation is set to 24 hours. The resulting output data<sup>1</sup> and scripts<sup>2</sup> can be found in public repositories.

### 4.2 SA-EA Algorithms

*SA-EA variants and the settings.* For this work, we are considering six different surrogates. We are considering *IDW*, *IDWR*, and *LR-KNN* as described in Section 2.2. *IDW* and *IDWR* do not require any settings, and are implemented as described in Section 2.2. For *LR-KNN*, the  $k$  is set to 32 to capture the local slope of the fitness landscape. In this work, we also consider simply using  $k$ -nearest neighbours (*KNN*), where the surrogate fitness is equivalent to the nearest solution in the set of solutions evaluated using the true fitness function. This surrogate is interesting to consider because this surrogate is fast to evaluate and does not interpolate between true evaluated solutions. Furthermore, many state-of-the-art SA-EA use Kriging. Therefore, this work also uses the Kriging surrogate (*Kriging*) from the well-known K-RVEA [8] as implemented in the PlatEMO [37]. Finally, a surrogate named *No structure* is included. This surrogate randomly selects a solution from the true evaluated solutions and returns the fitness values for *both* objectives. This surrogate functions as a baseline since it returns fitness values in the ranges of previous evaluated solutions, but there is no correlation between the true and surrogate fitness values.

All SA-EA variants are initialized with 32 solutions which are sampled using LHS. For all variants, the number of selected solutions per surrogate optimisation cycle,  $\mu$ , is set to 1 so that the surrogate is updated as frequently as possible (every new true evaluation). In this work, the MATLAB/Octave code implementation of RVEA from PlatEMO is utilized [37]. RVEA is initialized with a population size of 32 which is sampled using LHS. Furthermore,  $\alpha$  is set to  $1 \cdot 10^6$ ; this is intended to maximize convergence due to the low number of evaluations.

<sup>1</sup><https://zenodo.org/records/10575290>

<sup>2</sup><https://zenodo.org/records/10496995>



### 4.3 Dimensionality reduction

We seek to compare vectors of landscape features between true and surrogate fitness functions, and also between different surrogates. To this end, we use t-distributed stochastic neighbour embedding (t-SNE); this is a dimensionality reduction technique widely employed to visualize high-dimensional data in a lower-dimensional space while preserving the pairwise similarities between data points [39]. Employing dimensionality reduction individually for each sample within a time-series sequence (such as the different points in evolution which we track in this study) could introduce unwarranted variability in the consecutive projections. This complicates the process of tracking the optimisation. Therefore, also, a dynamic t-SNE [35] has been used in this work. For the experiments, we use a perplexity of 50, sigma optimisation iterations of 50, epoch count of 10000, and we set the random seed to 1. In the case of a dynamic t-SNE, the movement penalty is set to 0.001.

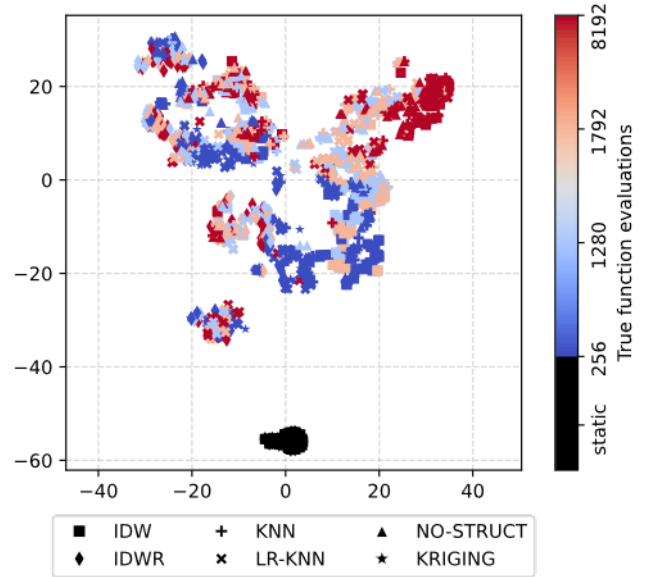
### 4.4 Performance Modelling

The modelling is conducted with random forest regression as the learning algorithm; it is used with its default hyperparameters, namely: 500 trees, with  $\frac{1}{3}N_f$  features included per split (where  $N_f$  is the number of features); sampling *with* replacement; the sample size is the same as the number of observations. The 28 features are computed using both the true fitness landscape and the surrogate, which leads to 56 features in the candidate pool. Four points in evolution are considered for the temporal approach: after 256 evaluations; after 1280; after 1782; and 8192.

For the static analysis, all features are based on the true fitness function and there is a single snapshot of the 28 features. The sample size for the temporal and static analysis is the same: 2000 (100D) points. The response variable is the median (normalised using the ideal and Nadir points) final hypervolume obtained by the surrogate-assisted RVEA variants. The IDW-assisted RVEA is excluded from the algorithm performance analysis: it achieved non-zero hypervolume on only one out of 55 functions, making it impossible to build a meaningful model.

We carry out recursive feature elimination (RFE)<sup>3</sup>, bootstrapped for 1000 iterations. In accordance with the one-in-ten guideline [15] for the ratio between features and observations, we limit the number of selected features to a maximum of five (there are 55 benchmark problems under study). Models are built using the features which were identified during RFE. To try and account for the potential effect of having a limited number of observations (the model could be susceptible to randomness in the data split), we bootstrap the models for 1000 iterations using a 80-20 training-validation split on the dataset. To quantify model performance, we consider the pseudo  $R^2$ ; this is computed as  $1 - \frac{MSE}{variance(y)}$  where  $y$  is the target variable (normalised final hypervolume achieved by the surrogate-assisted RVEA). If a model is very poor, the pseudo  $R^2$  can sometimes be less than zero; to preserve its meaning as the proportion of variance explained, we replace any negative values with zero. We report the bootstrap mean, median, and standard error of the model  $R^2$  on validation data.

<sup>3</sup><https://www.rdocumentation.org/packages/caret/versions/6.0-92/topics/rfe>



**Figure 2: t-SNE plot comparing the true features resulting from the temporal FLA to the static FLA. Each marker represents the median fitness landscape feature of a particular  $\text{BBOB-BIOBJ}$  problem. The colour of the point indicates in what phase of the evolution the landscape feature is calculated, as indicated with the colour bar.**

## 5 RESULTS

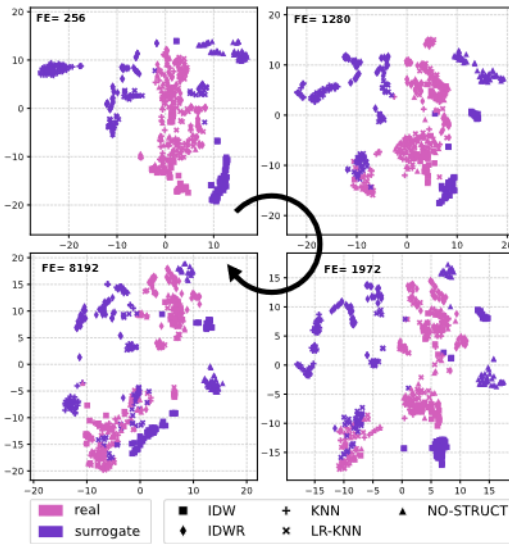
### 5.1 Static compared to temporal FLA

To ascertain whether the temporal FLA yields distinct features in comparison to a static FLA, we compared the true landscape features for both analyses in a t-SNE plot. The results are depicted in Figure 2. When using the *Kriging* surrogate, only 33 out of 55  $\text{BBOB-BIOBJ}$  surpassed 256 function evaluations before the optimisation runtime limit, see Table 1 in the supplementary material. Only these 33 problems are included in the further analysis. Each marker, in Figure 2, represents true fitness landscape features of one of the  $\text{BBOB-BIOBJ}$  problems. The black markers denote the features resulting from the static FLA, while the coloured markers indicate the phase of the evolution in the temporal FLA.

The most significant finding here is the distinct discrepancy between markers derived from static and temporal FLA, indicating that the two types of analysis exhibit divergent feature distributions. Moreover, the feature similarity between the different  $\text{BBOB-BIOBJ}$  problems appears to be larger than the similarity between the features at different stages of evolution (in particular, notice the red cluster in the top-left of the plot containing the majority of the late-in-evolution samples), emphasizing the need to perform fitness landscape analyses temporally. We have also evaluated, per feature, whether there is a statistical difference between the static and temporal analysis (see supplementary material).

### 5.2 True compared to surrogate landscapes

We also compare the true landscape features *and* the surrogate landscape features during different phases of evolution, see Figure 3. The top-left plot represents the features at the beginning of the



**Figure 3: Dynamic t-SNE plot comparing fitness landscape features for the true to the surrogate fitness evaluations acquired in the temporal FLA. Each marker in the scatter plot represents the median fitness landscape feature of a particular BBOB-BIOBJ problem**

optimisation (256 function evaluations). Following the plots clockwise, there are different phases in evolution, as indicated with the annotated text. In the plots, each marker represents fitness landscape features of one of the BBOB-BIOBJ problems. The colours now describe whether the features are based on the true or surrogate landscape. The surrogate (indicated using the marker shapes) is also relevant for the true landscape features, since different surrogates can lead the optimisation to different regions of the true fitness landscape. *Kriging* has been excluded for the analyses because we only have the data until 256 function evaluations. After 256 function evaluations, it is interesting that all true-landscape fitness features cluster together (notice the pink markers) while the surrogate features (these are coloured purple) disperse into the peripheral regions of the plot; this suggests a large difference between the true and the surrogate features. At the same time, the surrogate model data points are locally clustered together (into their surrogate types), suggesting there is also a clear difference between the features of different surrogates. To investigate the similarity between the surrogate and the underlying true landscape, we compare the corresponding feature distributions with a statistical test. Since in Figure 3, the different surrogates form individual clusters throughout the evolution, the statistical comparison is performed per surrogate – resulting in 588 tests. Because the surrogate landscape features are calculated on the same solutions as the true landscape features, a Wilcoxon test is used; we consider statistical significance to be  $p \leq 0.05$  (corrected using the Bonferroni correction).

Figure 4(a) summarises the results, where the black squares represent a significant difference in distributions. For *KNN*, 17 of the 28 surrogate features are significantly different from the true features throughout the evolution; for *IDWR*, *No structure*, *LR-KNN*, and *IDW* this is: 14, 10, 8, and 8 features respectively. Furthermore, *LR-KNN* seems to capture the true landscape from 1280 function evaluations:

this can be observed in the increase in light-grey squares when going up the vertical axis towards 8192 evaluations. This does not occur for *IDWR* and *KNN*. Interestingly, despite these discrepancies, when we consider the Spearman correlations between the true and surrogate landscape feature distributions, the features do seem to correlate with each other: see Figure 4(b). Overall, indicated by the number of blue squares, *No structure*, *LR-KNN*, and *KNN* seem to have higher feature correlation than *IDWR* and *IDW*. Note that, for *KNN*, there are many features whose distributions showed significant difference in 4(a) while simultaneously having high correlations. We also notice that several features appear to have perfect correlations between the true and surrogate distributions. In the case of *dist-x-avg*, *dist-x-avg-neig*, and *dist-x-max* – this is because these relate to distance in the variable space. For the others (*dist-f-avg*, *dist-f-max*, and *dist-f-avg-neig*): these three have highly skewed distributions for both true and surrogate landscapes, with the vast majority of values being low but with a few very large outliers. These distributions likely occurred due to the difference in fitness ranges across functions, and the nature of the distributions is probably the reason for the strong correlations. This raises the thought about whether this specific subset of features should perhaps be revisited or further developed in future research.

To understand the distributions further we also visualise the difference between the true-landscape median and the surrogate-landscape median of the (normalised) feature distributions: see Figure 4(c). For some features, particularly for *KNN* and *IDWR*, there is a large difference in the median of the feature distributions, explaining the significant differences seen in Figure 4(a).

### 5.3 Algorithm performance prediction

Table 1 presents a summary of the algorithm performance prediction models. Each row relates to a specific type of surrogate, and contains information about five models (one per column), where the response variable is the final hypervolume associated with using that surrogate. For example, the first row represents: four models using temporal features from sampling during runs using the *KNN*-surrogate, and one model utilising *static* features extracted from a latin hypercube sample. The type of sampling is indicated in the header row; for each row below that, *selected* reports the features included in the model (these were selected by RFE) – features of the surrogate landscape are in *blue text* and features of the true landscape are in *black text*; and beside the name of the surrogate we present the bootstrap mean, median, and standard error for the model pseudo  $R^2$  on validation data. The  $R^2$  values should not be taken as an indication of the performance of the surrogates, but as an indicator of how well landscape features can predict the outcome of optimisation.

Looking at the  $R^2$  values in Table 1 we can consider the values for the models built using a static landscape sample as the baseline for interpretation. For the *KNN* and *LRKNN* surrogates, three of the four temporal models are better than the static equivalent; the remaining one is the first in the time series (256 evaluations). For *IDW*, there are two superior temporal models, one which is roughly equal to the static equivalent, and one which is lower quality than the static counterpart (this is again the earliest temporal sample, at 256 evaluations). Finally, for the *No structure* surrogate, it is not




Figure 4: A) Comparing the true and surrogate landscape feature distributions yielded during different phases of evolution, where black refers to statistically significant ( $p \leq 0.05$ ) differences according to a Wilcoxon test. B) Correlations between the true and surrogate landscape feature distributions yielded during different phases of evolution. C) Difference in median between the true and surrogate landscape feature distributions yielded during different phases of evolution.

Table 1: Algorithm performance prediction models. The predictors are true (*black text*) and surrogate (*blue text*) landscape features sampled temporally at four points in evolution (and a static sample) chosen by feature selection; the response variable is final (normalised) hypervolume achieved by the surrogate-assisted RVEA variant indicated in the first column. Predictors and the response are the median over 15 runs. The model quality metric is pseudo  $R^2$ . Each model is bootstrapped for 1000 iterations and the reported values are *bootstrap mean* | *bootstrap median* (*bootstrap standard error*) for the metric on validation data.

sampling →	temporal (256 evals)	temporal (1280)	temporal (1792)	temporal (8192)	static
selected →	<i>[dist-f-max, dist-f-avg, dist-f-max, dist-f-avg-neig, diff-f-avg-neig]</i>	<i>[inf-avg-neig, slo-n, diff-f-avg-neig, dist-f-avg-neig, dist-f-avg]</i>	<i>[inf-avg-neig, plo-dist-max, dist-x-avg-neig, nd-n]</i>	<i>[diff-f-cor-neig, dist-f-cor-neig, lsupp-avg-neig, diff-f-cor-neig, rank-avg]</i>	<i>[dist-f-cor-neig, dist-f-avg-neig, dist-f-max, dist-f-avg, diff-f-avg-neig]</i>
KNN	$R^2: [0.755   0.753 (0.236)]$	$R^2: [0.819   0.855 (0.176)]$	$R^2: [0.913   0.858 (0.174)]$	$R^2: [0.947   0.934 (0.096)]$	$R^2: [0.697   0.785 (0.209)]$
selected →	<i>[sup-avg-neig, eval-aws, inc-avg-neig, dist-f-avg-neig, diff-f-avg-neig]</i>	<i>[diff-f-cor-neig, length-aws, sup-avg-neig, dist-x-cor-neig, dist-x-cor-neig]</i>	<i>[diff-f-cor-neig, sup-avg-neig, dist-x-cor-neig, dist-x-cor-neig, length-aws]</i>	<i>[slo-n, nd-n, dist-f-cor-neig, supp-n, diff-f-cor-neig]</i>	<i>[dist-f-cor-neig, sup-cor-neig, dist-f-avg, dist-f-avg-neig, diff-f-cor-neig]</i>
IDW	$R^2: [0.682   0.698 (0.223)]$	$R^2: [0.834   0.786 (0.201)]$	$R^2: [0.748   0.763 (0.204)]$	$R^2: [0.870   0.866 (0.181)]$	$R^2: [0.754   0.749 (0.192)]$
selected →	<i>[plo-n, rank-ent, plo-dist-max, plo-dist-max, dist-f-avg]</i>	<i>[plo-dist-max, slo-n, plo-n, plo-dist-max, slo-n]</i>	<i>[slo-n, sup-avg-neig, plo-dist-max, plo-n, slo-n]</i>	<i>[nd-n, plo-dist-max, slo-n, inf-avg-neig, slo-n]</i>	<i>[dist-f-cor-neig, dist-f-avg-neig, dist-f-avg, dist-f-avg-neig, dist-f-max]</i>
LR-KNN	$R^2: [0.867   0.832 (0.163)]$	$R^2: [0.967   0.837 (0.181)]$	$R^2: [0.945   0.846 (0.180)]$	$R^2: [0.923   0.836 (0.185)]$	$R^2: [0.871   0.754 (0.213)]$
selected →	<i>[diff-f-cor-neig, diff-f-avg-neig, dist-f-cor-neig, dist-f-avg-neig, dist-f-avg]</i>	<i>[slo-n, dist-x-avg-neig, dist-x-avg-neig, rank-ent, dist-f-cor-neig]</i>	<i>[dist-x-avg-neig, dist-x-avg-neig, dist-x-avg, dist-x-avg, slo-n]</i>	<i>[nd-n, nd-n, supp-n, supp-n, eval-aws]</i>	<i>[dist-f-cor-neig, dist-f-avg-neig, sup-cor-neig, dist-f-avg, diff-f-avg-neig]</i>
No structure	$R^2: [0.686   0.730 (0.215)]$	$R^2: [0.815   0.786 (0.171)]$	$R^2: [0.778   0.754 (0.202)]$	$R^2: [0.618   0.838 (0.172)]$	$R^2: [0.728   0.701 (0.223)]$

clear whether temporal or static analysis is the best approach: some of the temporal models are better than the static model and some are not. For both *IDW* and *No structure* surrogates, the models at 256 evaluations are worse than the static baseline. This is probably because the sampling, which is steered by surrogate fitnesses, is led to less relevant regions: this behaviour could be due to the limited archive of real-evaluated solutions at this point in the search and the poor ability of these two surrogates (indeed, they exhibit lower performance on the problems than *KNN* and *LRKNN*).

Notice from Table 1 that for each of the 16 models built using temporal features, a mixture of both surrogate and true landscape features were selected by the RFE. This hints that it is important to consider features of both the surrogate and the true landscape

when considering performance prediction for surrogate-assisted optimisation; however, to test this thought, we also built models using feature pools of 1) surrogate features only and 2) true features only. This experiment showed that generally, the surrogate-only and true-only models are of approximately equivalent performance to those in Table 1. That finding implies that, firstly, the surrogate models may be mimicking the true landscape well; and secondly, that features of either can be used in performance prediction. The results for these two setups can be found in Tables 2 and 3 of the supplemental material. In Table 1, it can be seen that fitness-based features (those containing *f* in either blue or black text) appear to play an integral role in many of the predictions: the majority of models has at least one included.

When it comes to these fitness-based features, more from the surrogate landscapes are selected when compared to the true landscapes. We notice that the type of feature which matters appears to change over the phases of evolution: when the 1792 timepoint is reached, there is a shift away from fitness-based features towards features which capture evolvability and ruggedness with respect to variable-space distances (those ending with *avg-neig* and *cor-neig*). Looking again at the temporal models overall, there is also a clear trend of features relating to local optima — both single-objective and Pareto — appearing in the models (for example: *slo-n*, *plo-n*, and *plo-dist-max*). Notice that this is particularly evident in the case of the *LRKNN* surrogate, and that these features seem to be important regardless of whether they are computed from the surrogate landscape or the true landscape. The prevalence of local optima features in the *LRKNN* models implies that the performance of this type of surrogate may be particularly affected by local optima.

In terms of model quality over time, for the *KNN* surrogate there is a trend over time: the model quality increases. There is no clear trend for the other three surrogates. As mentioned, the last column contains relating to the models using features of the static sample as predictors. We observe that the majority of selected features for these models relate to fitness in some way. Comparing the static models with the temporal equivalents, it can be seen that all four of the temporal *KNN* models out-perform the static *KNN* model; and three of the *LR-KNN* temporal models outperform the equivalent static model.

Things are less clear with the *IDW* and *No structure* surrogates. Taking into account the four temporal models, it could be said that the equivalent static models perform similarly to them. It may seem counter-intuitive that the *No structure* models have mean  $R^2$  values between 0.618 and 0.815. This surrogate simply assigns a random true-evaluated fitness to a solution. However, when we consider the details of these models, things become clearer. For example: in the case of the temporal (256 evaluations) model, the surrogate does not have access to many true-evaluated solutions yet so the sample for landscape analysis — and subsequently the predictive model — resembles that of the static sample in column 5. For the second and third temporal models, a lot of features based on distance in the variable space are selected (these begin with *dist-x*); this makes sense, because they are not directly related to the surrogate fitness function. The feature *dist-f-cor-neig* (surrogate) is selected for the 1280-evaluations model. This is a fitness-based feature (and therefore highly linked to the surrogate), but the fact it was selected for the models is understandable when we consider that this feature showed moderate correlation with its equivalent for the true landscape, as seen in Figure 4b). To conclude this section, we note that feature importances were computed for all models, and that generally, features within a model have similar importance to one another. The associated data can be found in Tables 3, 5, and 7 of the supplemental material.

## 5.4 Limitations and Discussion

The performance prediction results indicate that temporal fitness landscape analysis potentially be used for online surrogate selection in multi-objective optimisation. This ratifies what has been shown in a single-objective context [25]. We also found that while

many landscape features significantly differ between the surrogate and true landscapes, they are often correlated. A previous study [41] observed feature differences for single-objective problems. Additionally, our results showed that using either or both of the surrogate and true landscape features can be beneficial in performance prediction; to the best of our knowledge, this is a novel approach: previous studies have focused on the evolvability of the surrogate [25] or the fitness approximation error [3].

There are limitations to the conducted study. The feature selection approach — recursive feature elimination — can be affected by randomness in the data splits, and can sometimes fail to remove redundant features. Even so, every feature selection method has its own drawbacks: for example, evolutionary feature selection is computationally intense and may require parameter tuning. We also note that sampling strategies for conducting temporal landscape analysis introduce additional computational expense. Future work will consider how to best utilise the fitness evaluations which have already been carried out during optimisation. Another consideration is that the included surrogate models were chosen due to their straightforwardness; however, in future works, more Kriging models from other state-of-the-art SA-EAs such as in [8, 42] and other types of surrogate models such as neural networks in [13] should be considered. Finally, some of these models can become expensive to build and evaluate when the archive of solutions becomes too large. The impact of the fitness landscapes when reducing the archive size should also be considered.

## 6 CONCLUSIONS

Surrogate-assisted evolutionary algorithms can effectively exploit problem structures by substituting expensive-to-evaluate fitness functions. Analysing the resemblance and discrepancies between the surrogate landscape and the true landscape could provide a better understanding how surrogates can assist EAs in performing efficient optimisation by capturing various features of the true fitness landscape. In this work, we have investigated the difference between true and surrogate landscape features during evolutionary multi-objective optimisation. We consider several surrogate-assisted versions of a reference-vector guided evolutionary algorithm and use the well-known BBOB-BIOBJ suite of bi-objective functions. Our results indicate that surrogate landscape features differ significantly from the true landscape features and that these features vary during the course of a run. Despite these differences, the surrogate and true landscape often show a high correlation. This work also evaluates how these different features impact the actual search and identified key landscape features from both the surrogate and the true landscape with the capability to predict algorithm performance. This opens the door for online surrogate switching in multi-objective optimisation in the future.

## ACKNOWLEDGMENTS

This research is part of the research programme Open Technology Programme with project number 15586, which is financed by the Dutch Research Council (NWO), Elekta, and Xomnia. Further, the work is co-funded by the public-private partnership allowance for top consortia for knowledge and innovation (TKIs) from the Dutch Ministry of Economic Affairs.



## REFERENCES

- [1] Jason Adair, Gabriela Ochoa, and Katherine Malan. 2019. Local optima networks for continuous fitness landscapes. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 1407–1414.
- [2] Hanan Alsouly, Michael Kirley, and Mario Andrés Muñoz. 2023. Dynamic landscape analysis for constrained multi-objective optimization problems. In *Australasian Joint Conference on Artificial Intelligence*. Springer, 429–441.
- [3] Samineh Bagheri, Wolfgang Konen, and Thomas Bäck. 2016. Online selection of surrogate models for constrained black-box optimization. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 1–8.
- [4] Bernd Bischl, Olaf Mersmann, Heike Trautmann, and Mike Preuß. 2012. Algorithm selection based on exploratory landscape analysis and cost-sensitive learning. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*. 313–320.
- [5] Dimo Brockhoff, Anne Auger, Nikolaus Hansen, and Tea Tušar. 2022. Using well-understood single-objective functions in multi-objective black-box optimization test suites. *Evolutionary Computation* 30, 2 (2022), 165–193.
- [6] Gjorgjina Cenikj, Gasper Petelin, Carola Doerr, Peter Korošec, and Tome Eftimov. 2023. DynamoRep: Trajectory-based population dynamics for classification of black-box optimization problems. *arXiv preprint arXiv:2306.05438* (2023).
- [7] Ran Cheng, Yaochu Jin, Markus Olhofer, and Bernhard Sendhoff. 2016. A reference vector guided evolutionary algorithm for many-objective optimization. *IEEE Transactions on Evolutionary Computation* 20, 5 (2016), 773–791.
- [8] Tinkle Chugh, Yaochu Jin, Kaisa Miettinen, Jussi Hakanen, and Karthik Sindhya. 2016. A surrogate-assisted reference vector guided evolutionary algorithm for computationally expensive many-objective optimization. *IEEE Transactions on Evolutionary Computation* 22, 1 (2016), 129–142.
- [9] Kalyanmoy Deb, Ram Bhushan Agrawal, et al. 1995. Simulated binary crossover for continuous search space. *Complex Systems* 9, 2 (1995), 115–148.
- [10] Kalyanmoy Deb, Mayank Goyal, et al. 1996. A combined genetic adaptive search (GeneAS) for engineering design. *Computer Science and Informatics* 26 (1996), 30–45.
- [11] Kalyanmoy Deb and Himanshu Jain. 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Transactions on Evolutionary Computation* 18, 4 (2013), 577–601.
- [12] Leonardo Ramos Emmendorfer and Graçaliz Pereira Dimuro. 2020. A novel formulation for inverse distance weighting from weighted linear regression. In *International Conference on Computational Science*. Springer, 576–589.
- [13] Dan Guo, Xilu Wang, Kailai Gao, Yaochu Jin, Jinliang Ding, and Tianyou Chai. 2021. Evolutionary optimization of high-dimensional multiobjective and many-objective expensive problems assisted by a dropout neural network. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 52, 4 (2021), 2084–2097.
- [14] Nikolaus Hansen, Anne Auger, Raymond Ros, Olaf Mersmann, Tea Tušar, and Dimo Brockhoff. 2021. COCO: A Platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* 36 (2021), 114–144. Issue 1. <https://doi.org/10.1080/10556788.2020.1808977>
- [15] Frank Harrell Jr, Kerry Lee, Robert Califf, David Pryor, and Robert Rosati. 1984. Regression modelling strategies for improved prognostic prediction. *Statistics in Medicine* 3, 2 (1984), 143–152.
- [16] Kyle Robert Harrison. 2023. Surrogate-assisted analysis of the parameter configuration landscape for meta-heuristic optimisation. *Applied Soft Computing* 146 (2023). <https://doi.org/10.1016/j.asoc.2023.110705>
- [17] Chunlin He, Yong Zhang, Dunwei Gong, and Xinfang Ji. 2023. A review of surrogate-assisted evolutionary algorithms for expensive optimization problems. *Expert Systems with Applications* (2023). <https://doi.org/10.1016/j.eswa.2022.119495>
- [18] Nanda Dulal Jana, Jaya Sil, and Swagatam Das. 2018. Continuous fitness landscape analysis using a chaos-based random walk algorithm. *Soft Computing* 22 (2018), 921–948.
- [19] Anja Janković and Carola Doerr. 2019. Adaptive landscape analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2032–2035.
- [20] Anja Jankovic, Tome Eftimov, and Carola Doerr. 2021. Towards feature-based performance regression using trajectory data. In *Applications of Evolutionary Computation: 24th International Conference, EvoApplications 2021, Held as Part of EvoStar 2021, Virtual Event, April 7–9, 2021, Proceedings 24*. Springer, 601–617.
- [21] Anja Jankovic, Diederick Vermetten, Ana Kostovska, Jacob de Nobel, Tome Eftimov, and Carola Doerr. 2022. Trajectory-based algorithm selection with warm-starting. In *2022 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 1–8.
- [22] Ana Kostovska, Anja Jankovic, Diederick Vermetten, Jacob de Nobel, Hao Wang, Tome Eftimov, and Carola Doerr. 2022. Per-run algorithm selection with warm-starting using trajectory-based features. In *International Conference on Parallel Problem Solving from Nature*. Springer, 46–60.
- [23] Ryan Lang and Andries Engelbrecht. 2019. On the robustness of random walks for fitness landscape analysis. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 1898–1906.
- [24] Ryan Dieter Lang and Andries Petrus Engelbrecht. 2020. Distributed random walks for fitness landscape analysis. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. 612–619.
- [25] Minh Nghia Le, Yew Soon Ong, Stefan Menzel, Yaochu Jin, and Bernhard Sendhoff. 2013. Evolution by adapting surrogates. *Evolutionary Computation* 21, 2 (2013), 313–340.
- [26] Jin Li and Andrew D Heap. 2011. A review of comparative studies of spatial interpolation methods in environmental sciences: Performance and impact factors. *Ecological Informatics* 6, 3–4 (2011), 228–241.
- [27] Jian-Yu Li, Zhi-Hui Zhan, and Jun Zhang. 2022. Evolutionary computation for expensive optimization: A survey. *Machine Intelligence Research* 19, 1 (2022), 3–23.
- [28] Arnaud Liefvooghe, Bilel Derbel, Sébastien Verel, Manuel López-Ibáñez, Hernán Aguirre, and Kiyoshi Tanaka. 2018. On Pareto local optimal solutions networks. In *Parallel Problem Solving from Nature—PPSN XV: 15th International Conference, Coimbra, Portugal, September 8–12, 2018, Proceedings, Part II 15*. Springer, 232–244.
- [29] Arnaud Liefvooghe, Sébastien Verel, Benjamin Lacroix, Alexandru-Ciprian Zăvoianu, and John McCall. 2021. Landscape features and automated algorithm selection for multi-objective interpolated continuous optimisation problems. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 421–429.
- [30] Katherine Malan and Andries Engelbrecht. 2013. A survey of techniques for characterising fitness landscapes and some possible ways forward. *Information Sciences* 241 (2013), 148–163.
- [31] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. 2011. Exploratory landscape analysis. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*. 829–836.
- [32] Mario A Muñoz, Michael Kirley, and Saman K Halgamuge. 2014. Exploratory landscape analysis of continuous space optimization problems using information content. *IEEE Transactions on Evolutionary Computation* 19, 1 (2014), 74–87.
- [33] Zbyněk Pitra, Jan Koza, Jiří Tumpach, and Martin Holeňa. 2022. Landscape analysis for surrogate models in the evolutionary black-box context. *arXiv preprint arXiv:2203.11315* (2022).
- [34] Zbyněk Pitra, Jakub Repický, and Martin Holeňa. 2019. Landscape analysis of Gaussian process surrogates for the covariance matrix adaptation evolution strategy. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 691–699.
- [35] Paulo Rauber, Alexandre Falcao, Alexandru Telea, et al. 2016. Visualizing time-dependent data using dynamic t-SNE. (2016).
- [36] Donald Shepard. 1968. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM National Conference*. 517–524.
- [37] Ye Tian, Ran Cheng, Xingyi Zhang, and Yaochu Jin. 2017. PlatEMO: A MATLAB platform for evolutionary multi-objective optimization. *IEEE Computational Intelligence Magazine* 12, 4 (2017), 73–87.
- [38] Marco Tomassini. 2022. A Local optima network view of real function fitness landscapes. *Entropy* 24, 5 (2022), 703.
- [39] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [40] Diederick Vermetten, Hao Wang, Kevin Sim, and Emma Hart. 2023. To switch or not to switch: Predicting the benefit of switching between algorithms based on trajectory features. In *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*. Springer, 335–350.
- [41] Bernhard Werth, Erik Pitzer, and Michael Affenzeller. 2020. Surrogate-assisted fitness landscape analysis for computationally expensive optimization. In *Computer Aided Systems Theory—EUROCAST 2019: 17th International Conference, Las Palmas de Gran Canaria, Spain, February 17–22, 2019, Revised Selected Papers, Part I 17*. Springer, 247–254.
- [42] Qingfu Zhang, Wudong Liu, Edward Tsang, and Botond Virginas. 2009. Expensive multiobjective optimization by MOEA/D with Gaussian process model. *IEEE Transactions on Evolutionary Computation* 14, 3 (2009), 456–474.